

[SQUEAKING] [RUSTLING] [CLICKING]

MICHAEL SCOTT A. CUTHBERT: OK, we're going to go through a little bit of a topic that you've unlocked now with the end of problem set 3, which wasn't totally-- wasn't officially locked, but we didn't really talk about it. And that's stream filtering. So let me put this down, and let's do a little bit of Jupyter work. So go ahead and get yourself situated in that. I'll give you a second or two. Uh-oh.

[SIDE CONVERSATION]

And I think I saw scores generated in text for everybody. There was a great question on the list. What do you do if you're not yet on version 9, which none of us are, and you want to see more of the score, you can do `show musicxml dot pdf`. And if you have MuseScore, that will automatically generate your PDF of more than just the first page. You can also just do `show musicxml`. So those should be straight quotes. I'm just used to putting them as curly quotes. You all know that, right? And that will just open it in your favorite notation editor.

Great. So let's get started with-- let's not load our favorite Bach chorale, but let's load our new favorite piece by a Boston-based composer, Amy Beach. And it's a piece for-- you see I even start typing Bach. It makes life-- oops. Aah! When I'm not teaching in the semester I have-- well, it's import star all right now, so we don't have to remember everything. But not every module is unlocked yet. Usually, when I'm not teaching class, I make it so that that's automatically imported for me so I don't have to type that. So if you see me make that mistake, that's what it is.

We'll just start by looking at this piece. And it has a similar structure to the piece that we have been doing. Ooh, we have some layout errors today on my thing, but it's a piece called *Prayer of a Tired Child For four Part Course of Women's Voices*. And see that the four parts are soprano 1, soprano 2, alto 1, alto 2. And then there's a fifth part, the piano. And the piano part has-- well, we'll see what it is in just a bit.

So how would we get the soprano 2 part? What do we want to do? You all have gotten pretty good at manual filtering. So beach dot something. Actually, talk with your neighbors a little bit about that. We have some odd number neighbors and try to talk through what you do. It's more important that you can explain things than do it instantly.

[SIDE CONVERSATIONS]

MICHAEL SCOTT A. CUTHBERT: You want to get the soprano 2 part.

[SIDE CONVERSATIONS]

MICHAEL SCOTT A. CUTHBERT: So good, good. One way. If you find one way, keep going for another or two until we get it. There's about five different ways.

[SIDE CONVERSATIONS]

MICHAEL SCOTT A. CUTHBERT: Not good here. Everybody, give a thumbs up if you think how to do it. If not, keep talking with-- because that's almost universal thumbs up. Jonathan, was that a thumbs up? Yep, good. OK, great. Let's see. Jordan, how would you do it?

AUDIENCE:

Dot parts and brackets for 1.

MICHAEL SCOTT A. CUTHBERT: Brackets for 1. Good. So we've already started with one of the things on today. This was in one of the users guide thing. Dot parts is a kind of filter that filters out everything that is not a part object. And so this will let you find things. Good. What's another way we can do it? Somebody else. Who have I not been calling on so much? Derek.

AUDIENCE:

You can put in quotes the-- as the index.

MICHAEL SCOTT A. CUTHBERT: For the index to parts. Oops. Sometimes that autocomplete is less than helpful. Oops! What do I have to do? What am I looking for? Somebody help? Soprano 2. Yep. Great. So it has to be exact. Does not grab the first one. That kind of makes that-- good. What's another way? Somebody else? Vanessa?

AUDIENCE:

Is it by ID?

AUDIENCE:

I did the first one.

MICHAEL SCOTT A. CUTHBERT: You did the first one? Great. Matthew?

AUDIENCE:

getElementByld?

MICHAEL SCOTT A. CUTHBERT: Good. getElementByld Soprano II. That's helpful if you don't know, well, for instance-- well, not piano. Let's try that dot getElementByld Piano. You actually get nothing in that case. Though, we'll see that there's some interesting things when it comes to piano. Let's do this for p in beach dot parts, print p. And we can actually see a couple of things here. One is that the piano is represented by 2 or something called a part staff that are separate parts in the ontology of music²¹.

In a lot of other systems, they are the same part, and there are two staves within a part. And every one of these representation things has huge drawbacks to them. And when somebody comes up with a representation for piano music that does not have the drawback, you can please, please publish that and you will make the world so happy. Because everything is fine with treating a piano as two parts. Let's say it's in 3 over 4. As two separate parts as long as-- I'm trying to make this at least vaguely consonant-- as long as you do something like this.

It's when you start doing-- that should go up to do proper stemming. That you have the problem, right? That suddenly, a musical melody that starts in one part continues into another part. Does your ontology allow for beams to cross parts?

Well, I never-- you asked me to create an ontology for something that wasn't Western classical music, I have no ontology for that. Yeah, but this ends up being one of the main things. Whenever somebody says they've created a new way of representing classical music, the thing I always do, because I'm a smart aleck, and if you want to be a smart aleck, just always go up to that person who's like, oh, I have a much simpler way. Just ask them, How do you do cross-staff beams in your ontology?

And if they go, well, we haven't gotten to that, this is the first thing that almost every ontology should start with. How do we represent cross-staff beams? If you figured out how to represent cross-staff beams or chords that span-- what would be a good chord on this? I guess we're in first inversion here. Chords that span multiple staves, you'll figure it out. The music21 ontology on this is bloop. It's really not what it should be.

On the other hand, when you represent everything as one staff and you have to say, OK, well, this is a note on staff 1, this is a note on staff 2, then your left to right parser ends up failing quite a bit because you go left to right. And then you say, OK, now let's jump backwards a bit. And when you have some sort of array, some sort of a sequence of objects that go all in a row, but then sometimes jump backwards and fill things in, I guarantee you, you will be managing what we call cursor hell for the rest of your life, trying to figure out where you are left to right and how, if you're thinking about your big O algorithmic notation, how can we simplify things? How can we make walking through a score be O of n? Or how can we do-- if we guess a number between 1 and 100, I tell you higher or lower, what's the big O we're talking about?

Log n.

AUDIENCE:

MICHAEL SCOTT A. CUTHBERT: Log n. Good, good. Yeah. Some things MIT does really, really well with, and this is great. A lot of things I might use as well. Yeah, so you all know log n. Well, what happens to your log n guessing if you have an array of things that go 0, 1, 2, 5, 11, 14, 8, 2? Suddenly, what are we back to, to find any element in something? O of n. You have to look at every element. So you can't know where something is. Yes, Adam.

AUDIENCE:

You were reciting the Catalan numbers for a bit.

MICHAEL SCOTT A. CUTHBERT: I think I was, yes. But something we will get to later in the class, the strong law of small numbers. You can look this up somewhere. It is the numbers-- numbers. It is one of the biggest mistakes with computational work in any non-traditionally computational field like musicology. And that is to say-- you don't have to look it up now. You can look at it later. Basically, there aren't enough small numbers for all the things that you want. So anytime you're giving out small numbers, there's almost always going to be a pattern. It's going to be Catalan numbers. It's going to be Fibonacci numbers. And people always find patterns in all these things.

OK, that was a digression on a digression on a slight digression. So we'll go back to the main thing, which is finding parts in here. Good. So we have a way of going through and finding interesting parts. I'm going to turn this off so we can see better. Is that a little bit better? Yeah. Good. OK. Great. So we can figure out all the parts. Some of the other things, you can do it for for thing in beach dot recurse if isinstance-- help me. Let me start with this if isinstance. Be my GitHub autopilot. I've checked this. Oh good, good. Some people have played with this. If isinstance--

AUDIENCE:

Thing.

MICHAEL SCOTT A. CUTHBERT: --thing comma. What are we looking for?

AUDIENCE:

Stream dot part.

MICHAEL SCOTT A. CUTHBERT: Stream dot part if-- well, we can say here, we can start with this part_num equals 0 to 1 to negative 1. We'll say negative 1. part_num equals 1. Here, we'll say Soprano II equals none. part_num not equals 1, plus equals 1 if part_num equals-- wait. So we start with negative 1. So we get to 1. Soprano II equals thing. And let's see if that works. Coding on the fly. OK good, I can still do that. That's a really, really brutally awful way of doing it, right? But you can do it that way.

So there's one other way that I'm going to get you. And why did I come back to the recurse way of doing this? Because there is one particular shortcut that you're going to find very useful. And that is-- so beach index, you know this one, right? Index 0. Oh, this one, for those of you who just assumed that everything starts with a lot of things, we have a whole bunch of text boxes. What happens around 10? We're still doing text boxes, which appear somewhere on-- oh, there we go.

Part 19 is Soprano II. But here's another way that we can start getting to parts. This is a little bit different we have. But instead of putting a number in there, we can put in a class. And that lets us go through everything. So we can say for p in beach stream dot Part, print p. And that will be one of your simplest ways of going through things. For instance, for r in stream of beach-- sorry, I should scroll down a bit, beach notes dot rest print r. And now, you can quickly get a sense of what durations are used in every rest in the piece. Sorry, I'm not going to scroll down because people want to see that.

Give you a chance to catch up and see, What are some of your favorite things? Does the piece have any clef changes? Or actually, let's just not say clef change. Does it have any time signature changes? So let's see how quickly you can have the computer figure out, does this piece, change time signature. Have any-- no. Are there any clef changes? Let's try that. It's all code, one behind you. All code doesn't have any time signature changes. In case you don't have this. What do we think about clef changes? Yes, no. No. We got yes. We got yes. We got yes, it does have some clef changes.

So shall we do the same-- I'll let people have a little bit more, and then we'll let a yes give me the code for doing this. OK, I finished talking. Who's a yes? Jason, you were you yes or-- yeah. So what should I be typing?

AUDIENCE:

for thing in beach dot recurse--

MICHAEL SCOTT A. CUTHBERT: Beach recurse.

AUDIENCE:

--if isinstance.

MICHAEL SCOTT A. CUTHBERT: OK, let's go to the new thing.

Instead of if isinstance, what do you want to do instead of that? You're on the right track, but now we have a new way of doing it. So this is a shortcut for recurse and things.

Jonathan.

AUDIENCE:

clef dot Clef.

MICHAEL SCOTT A. CUTHBERT: Clef dot Clef. Yeah. So it's basically the same thing. And, oh, here is something that came up in a couple of things. Now I know it's going to be a clef. So maybe thing isn't the best. So do you like that? No, this comes up a lot. What-- Adam doesn't like it. Somebody else? Anybody else not like this? Yeah, Paul, what do you not like about it?

AUDIENCE:

Because clef is part of the music21 input.

MICHAEL SCOTT A. CUTHBERT: Yeah, because clef is part of the music21 input. Some people who hate that I recommend importing all of music21, so you could do this if you wanted. And then clef is only referenced with-- what do you call it? With a prefix of music21. And some people really insist on that. You might have a friend who talks about don't pollute your namespace by things you don't know, which can be pretty good advice. But, yeah, you can do that. So thing is not bad. I sometimes do leave off the vowel. Some people don't like that.

Somebody in P set 3 did a lot of this with no putting an underscore afterwards. That's a pretty good thing. The other thing I will often do is just use the first letter, even though I usually like verbose things if I'm going to do that. So OK. And Jason or Jonathan, whoever was starting wants to go next. Print c, is that what you did, or? Yep. Right. I'm going through well, and we have a clef change. So clef change. Anybody object? Angelica? [MUMBLES] Adam has a bit-- yeah, Vincent.

AUDIENCE:

I tried that at first, but I really know if the bass clef maybe belonged to the piano part, piano change. And so maybe you can probably print out the part and then the clef that it corresponds to.

MICHAEL SCOTT A. CUTHBERT: Oh, OK.

AUDIENCE:

It is just in the part.

MICHAEL SCOTT A. CUTHBERT: Good. So what we could do, let's see. How can we get the part name on this? For-- yeah, go ahead.

AUDIENCE:

I just have another form for that that leads to the parts bit.

MICHAEL SCOTT A. CUTHBERT: For p in beach. And now we'll-- so the difference by the way, dot parts is not recursive. But hopefully, the parts are at the top most level of the score. So it doesn't matter. We could do either way, yep, and then for print p, maybe. Print p. Actually, I usually like to do something like this when I have an outer loop. Oops. Then we'll do for c in beach clef dot Clef print c. What did I do wrong?

AUDIENCE: It recurses.

MICHAEL SCOTT A. CUTHBERT: A recursing and beach. What should I be doing?

AUDIENCE: p.

MICHAEL SCOTT A. CUTHBERT: p clef dot Clef. Good. In any life, don't let me make a fool of myself, especially not when we're filming this class. So yeah, call me out when I do something like that, but great. So now we can see, ah. Oh. So there actually is a clef change, but not in one of the voices I was thinking. None of the voices did, but the piano part does change clef. Here's another way we can this. I'm going to be unlocking some nice stream things for this. We can just skip that outer loop.

For c-- you can just keep going and doing things-- in beach clef dot Clef. Let's do this. Print c comma. I'll let everybody get caught up on that. c dot getContextByClass stream dot part, and make sure to close your second parentheses. So this is a way of moving upward in a hierarchy to try to get what the most recent part is. Same type of information, but maybe a little bit easier for your computer to parse. What should we do if we want to know? Let's see where that bass clef change takes place. Let's assume it's like Wagner's *Lohengrin*, and we don't want to switch the whole thing. Yeah, Misha.

AUDIENCE: I can get contracts by class for a string, not measure.

MICHAEL SCOTT A. CUTHBERT: Great. So I'm going to cut and paste because we don't want to spend this whole class watching me make typos. Context by class, stream dot Measure. Whoa, it's a lot of things. When you're zoomed in, you don't get very much. Maybe I should print that in separate lines, but I want to look at the bottom one. Measure 29 offset 112. Here's a trick that I explained at the beginning of this class. I'll open it up in a PDF reader. And hopefully, this one has measuring numbers defined. Yeah, 112. No, 29. Phew! 29. There we go. Finally, there. What is that-- give me some ideas. What's the piece scored for? What are the instruments in this?

AUDIENCE: Women.

MICHAEL SCOTT A. CUTHBERT: Women's chorus and--

AUDIENCE: Piano.

MICHAEL SCOTT A. CUTHBERT: Piano. What are you going to guess, unless Amy Beach likes to use 20 million ledger lines. There's something about the piece. About the piano part. That's right. Misha.

AUDIENCE: It's pretty high up relative to--

MICHAEL SCOTT A. CUTHBERT: Yeah. So it's up with the voices the whole time until right at the very end. And so you can probably take a guess that something dramatically changes at the end. And it doesn't just go down because, well, we're tiny bit-- where's that note? We're tiny bit too low for treble clef. So we're going to drop down into bass clef. No. No, we get very, very low at the beginning.

So you could think about it. Immediately, you could ask, What are pieces that have a dramatic change in range only at the very end? Anytime you see something like this, always be thinking of, What are the next three questions you can go from here? So this is pretty interesting type of research problem that becomes faster to find with the computer. We could also-- you all know how to find the highest or the lowest note in a part, right? Yeah. OK. You're all experts at that now after problem set 3. So we can figure out this thing from here.

By the way, given the-- let's look at the beginning and stuff. And is there anything-- always take a second when you have a new piece. I should-- there isn't a good recording of this piece. If anybody sings in a four-part women's choir or a cappella group that could do it, it would be really nice to get a good recording of this. But where is the interest in this piece? Is it in soprano 1, soprano 2, alto 1, or alto 2? Or more than one place. Adam, so we'll give other people just a little bit of time. Adam, where do you think interest is?

AUDIENCE:

The middle two.

MICHAEL SCOTT A. CUTHBERT: The middle two. We asked this before. Who sings a middle voice, an alto, a mezzo soprano, tenor voice? Some people. How often is that the most interesting part in the piece?

AUDIENCE:

Almost never.

MICHAEL SCOTT A. CUTHBERT: Almost never. OK, so either Amy Beach is very, very kind to middle-voiced people, or there might be another reason. What's the name of this piece? I've told you a little bit ago. You have it on your computer. See how much that sticks in when we're doing these kind of corpus analysis. I'm going to give it back to you in a second. When we're given one of these corpus analyzes, try to think of what kinds of questions you might answer based on thing. *Prayer of a Tired Child*. How do you think-- I don't know why the layout is so bad today, but I think there might be any connection of the title to this opening? Yeah. I want to just spread it around it, yes. Well, first, what makes it a prayer? Yeah.

AUDIENCE:

Starting with "Our Father."

MICHAEL SCOTT A. CUTHBERT: Starting with "Our Father," yeah. Depending on your religious tradition, for some Christians, that can be a pretty standard thing. What makes it, do you think, *Prayer of a Tired Child*?

AUDIENCE:

I'm trying-- I'm not sure.

MICHAEL SCOTT A. CUTHBERT: OK. Good. Call on somebody else. Let's put somebody else on the spot. Vanessa, what makes it a tired?

AUDIENCE:

I mean, she does, in fact, say here the

MICHAEL SCOTT A. CUTHBERT: Title. OK, so the text.
Definitely good. Good. Always go with that. Yeah, Angio.

AUDIENCE: I think it's how soft it starts out.

MICHAEL SCOTT A. CUTHBERT: So it's definitely soft. Yeah,
so we want to do that. Good. Anything else? Maybe
Jonathan.

AUDIENCE: I mean to the part of distance. That seems to say no over and over.

MICHAEL SCOTT A. CUTHBERT: Right, so what might that
have to do with tiredness?

AUDIENCE: I mean, you're just not moving.

MICHAEL SCOTT A. CUTHBERT: You're just not moving. Yeah.
You're either not moving because you're falling asleep, or
you're not moving because you don't have the energy to
move up or down. So this is one of the things you can start
thinking about, how in a text can-- how might the text have
something to do with the amount of motion in a piece? So
we can look at that. And then finally, finally, you measure
for-- they finally get to move. And which direction does the
soprano 1 move when it finally gets to move?

AUDIENCE: Down?

MICHAEL SCOTT A. CUTHBERT: Down. OK. So these are
some of the things that you guys were putting in your
questions on the Bach chorales that can apply to other
pieces, not just by Bach. By the way, how long-- yeah. Then
when the lowest voice, I always want to say bass, but alto 2
moves, where does it finally go? Down also. Great. So some
of the things we can do with this. Now, let's get back to
filtering back to the piece. Let's say we want to see this
measure here. Let's grab the entire context for where this
happens and just show it. So what we can do is we can do
beach dot measures, let's say 28 to 30, dot show. Hopefully,
I got the right one. Yep.

So, whoa, everything gets cleaned up when I don't show the beginning. So you can get-- another one of these commonly used filters is the measures thing. Now, I don't know if this was a mistake in ontology or not, but I've decided that measures will be based on how a musician thinks of them. So 28 to 30 is measure number 28 to measure number 30 inclusive. So some of you guys saw this. What are some things that can cause the number of, let's say, bar lines minus 1 to be not the same as the measure number, the highest measure? What are some things that might look like measures that don't get counted as measure numbers in your experience? Yeah, Vanessa?

AUDIENCE:

Pickups.

MICHAEL SCOTT A. CUTHBERT: Pickups. Great. Pickups. So did you have another one that you're saying? No. Good. Anybody else think of something?

AUDIENCE:

If you have a Coda, you do that.

MICHAEL SCOTT A. CUTHBERT: Ah, if you have a Coda or a repeat, and then you expand the repeat, we tend to call that measure 1, 2, 3, 4, 5. Well, we'll say 6 because you usually repeat in even numbered measures. And then you repeat it. We'll sometimes call that like 1a, 2a, 3a, 4a, 5a, 6a. The measure after 6 is usually 7. It is not 13 if you repeat, right? We don't count that. Here's another commonly encountered one. Especially, it seems to happen in 3 4. Here. Sorry, let me start that again. Pick up measure.

[SCAT SINGING]

We want to see this, et cetera. Where you have an incomplete measure or repeat sign, and then you begin the pickup again afterwards. Generally, we'd call this measure 1, 2, 3, 4, 5. But most computer representations will represent this as a measure and this as a measure that are both part of the same conceptual measure. Why? It's easier sometimes to only look for repeat signs at the ends of measures, and not have to look between every note to see if there's a repeat sign. So those are some of the reasons that might call off. So here's how we can do it this way.

Another thing that you'll find quite often, and it said what the offset of that-- offset 112, we could say, well, piano left hand equals beach dot parts, the final part. And then piano left hand getElementsByOffset. What do we say? Well, let's grab a little bit of context, 100 to 150. Oops. And we want to cast that into a list or something like that.

And so that lets you get things from within a smaller area. You'll find there is also-- because getElementsByOffset is not recursive. There is a recursive version, getElementsByOffset in hierarchy. And I wish I had come up with a shorter name for that, but-- so you can be able to find anything at a particular element within a hierarchy.

So these are some standard things. Some other people were putting in using lambda filters in some of the code do that. Who has worked with a lambda filter before in their life? Not every-- but most people. OK. So you can write a little piece of code to say, just get me things that are on here. One reason you might want to use some of these built-in features is that it's a little bit expensive to iterate over a stream. So the filters do sort of a pre-iteration.

For those who are in finance, it's kind of like a tax-exempt savings plan. You only get out the elements and unwrap them and make them look really nice once you have them. Peek under here, see if there's anything else I wanted to do. Nope, that's basically it on-- any other questions? Yeah.

AUDIENCE:

Yeah. So during the piece, I figured like we covered this at some point in class, but when you flatten, it loses some of the information when I try to then get context by class. Like, for example I think, yeah, it assumes everything is the last part even though it's not. Is that because of how flattening there is?

MICHAEL SCOTT A. CUTHBERT: Yeah, it's because it's-- here, so let's do this. We'll say beach flat equals beach dot flatten. And now here, we'll look at the first like 100 things in beach flat. And yeah, so we get everything that's here. And within that, so we're getting all these things here. Almost all the other get contexts will work really well. We'll get to slurs and dynamics and how they're represented because they're fascinating later.

But the thing is beach flat itself is a score with, in this case-- sorry. So in some cases, if you flatten a part, it becomes a part. If you flatten a score, it becomes a score. So if you flatten the part, you won't be able to get its original part context, but you should be able to get it with this. Were you not able to--

AUDIENCE:

Yeah, like if you iterate through these and then-- through the notes in the flat, for example, will you be able to get the correct-- all the parts?

MICHAEL SCOTT A. CUTHBERT: Yep, let's see--

AUDIENCE:

If they should be different.

MICHAEL SCOTT A. CUTHBERT: Let's hope so. Class stream dot part. Sorry, it's at the very end.

AUDIENCE:

I think that's--

MICHAEL SCOTT A. CUTHBERT: Oh, interesting. That seems to be a bug. I'll look into that. See what might have caused that to come in. We have a pretty rigorous testing, but I don't remember ever testing on that particular one.

AUDIENCE:

That's what we have here.

MICHAEL SCOTT A. CUTHBERT: Yep, that's why I have you. Thank you. Yep. Yeah, that'll be interesting. What you might be able to do, I'm not going to try to-- no part of this class is about working around music21 bugs, though you will find them. Let's see what happens if we do get context by class stream dot measure. OK. OK, so that one does change. So I suppose you could do this. You get the streams measures context and then get that measures context by class part.

And now we have different stuff. I'll try to figure out exactly what's happening to not make that work properly. But it was happening with the Bach chorales also, right? So it's not a piano bug. The piano parsing still has bugs. That's why you're not going to be working too often with piano scores. Only when I've really done that. Great. Any other things before we move on to the next and most fun topic for a little bit? OK, I'm going to move over to the board for a little bit. And let's just do some by hand warm-ups. We're going to be moving on to intervals for a bit.

And these will all-- supposedly notes that are supposed to be lined up. Just write them down here. I'll give them as number 1, 2, 3, 4, 5, 6, 7. Why don't you just say-- and actually, I won't put them as chords. I'll always put them here so that there's a directionality on them. And go ahead and work through as I'm writing them what the interval is in terms of not number of semitones, but we'll call it-- sometimes we'll call this the diatonic interval. These are one of these things people haven't completely done. So like major third, perfect fourth, things like that for all of these. Just do it as a review. Second one is always right.

Sorry. Sorry, I'm going to change number 4 for a second. I want to have at least one. Goes the other direction. I'll put this up for a second. OK, just so we're on the same page. Everybody, just one second. Everybody shout out number 1 on the count of 3. 1, 2, 3.

Major third.

AUDIENCE:

MICHAEL SCOTT A. CUTHBERT: Major third. OK, good. Keep working. I just want to make sure we're all using the same system. So let's get back together. Even if you don't have them all, we have a major third at the top. And we'll start with Angelica. And we'll move-- no, we'll move this way. What did you get for this one? Do you have that?

Diminished fifth.

AUDIENCE:

MICHAEL SCOTT A. CUTHBERT: Diminished fifth. Anybody get another answer? It's good. Anyone? Any other answers? Yeah.

Diminished 12th.

AUDIENCE:

MICHAEL SCOTT A. CUTHBERT: Diminished 12th. Ah. Any other answers? Augmented fourth, anybody? We have augmented fourth on the table. OK. So we have a couple of different contexts where any of these things can be right. Oh, and on this one also, I should have been pretty anal from the beginning. If I want to say a direction before the major third, what do we want to say? Up or let's be a little bit more elevated. It is an--

Ascending.

AUDIENCE:

MICHAEL SCOTT A. CUTHBERT: --ascending major third.
Good. This is an ascending. I'm just going to do that for a little bit, so we can keep going with that. So it's a diminished fifth. What's a context where diminished fifth is as good an answer as diminished 12th and what's one that doesn't? What are we assuming? Which type? You've been doing some of the reading. What's the equivalence class that we're using that allows this?

AUDIENCE:

Octave equivalence.

MICHAEL SCOTT A. CUTHBERT: Octave equivalence. So the O of OPEC, if we allow the O equivalence class, then we don't care if it's diminished fifth or diminished 12th. When might-- who is a stakeholder who might care about, Can you sing this-- can you sing this piece? It has a whole bunch of diminished fifths in it. I go, yeah, no problem. I'm an experienced singer. Great. So you're going to sing-- what note's that? You're going to sing this note, and then you're going to sing this note. You said you had no problem singing diminished fifths, right?

So there are contexts where we will very much care about not having octave equivalents. And there's other ones where we will. I'll give you the music21 terms for how these things work. But we sometimes call this a simple versus a compound interval. So sometimes we care only about simple intervals. Sometimes we care about compound. When's a time that we might be OK with calling it an augmented fourth? What equivalence class might be in play?

AUDIENCE:

Spelling?

MICHAEL SCOTT A. CUTHBERT: Spelling or enharmonic? Yeah. So if spelling doesn't matter, what's an app that you might make-- I have made, I suppose-- where when you're doing this kind of interval identification, you want to be careful and turn off spelling for a bit? What's a context of interval identification, where you don't want to be too uptight about whether they hear it as-- whether I gave it away kind of diminished fifth or augmented fourth? Yeah.

AUDIENCE:

It's the one where you're just trying to train your interval training.

MICHAEL SCOTT A. CUTHBERT: Yeah, where you're doing ear training, right? You can't usually hear the difference between a diminished fourth and augmented fifth, at least not till you hear the next note, right? When you hear the opening of *The Simpsons*, then you know it's augmented fourth and not a diminished fifth based on where it goes. Good. OK, then next one. Let's keep going. I don't want to spend too much time on this. I want to get some programming. Paul, do you want to take a shot on this one? You can take a shot or pass. Go ahead. But you take a shot.

AUDIENCE: Augmented second.

MICHAEL SCOTT A. CUTHBERT: Augmented second. And just to be pedantic, what direction?

AUDIENCE: Ascending.

MICHAEL SCOTT A. CUTHBERT: Ascending augmented second. Good. Great. I don't think there's any way that you could say it's an augmented ninth. You can always reduce the octave, but you cannot-- it's an interesting non equivalence class, right? You would never be allowed to call that augmented ninth. But you can call this one a diminished fifth. Great. Number 4. Let's continue, Jonathan.

AUDIENCE: Descending augmented octave.

MICHAEL SCOTT A. CUTHBERT: Good. Descending augmented and augmented octave. Great. Now let's keep going around to Matthew. What do you want to call number 5?

AUDIENCE: For lack of a better word, I might say augmented augmented augmented second?

MICHAEL SCOTT A. CUTHBERT: Augmented augmented augmented second. Everyone agree with that? Well, let's look at this. So D to E is what kind? What's that note? What's that interval-- major second. So D to E sharp is?

AUDIENCE: Augmented.

MICHAEL SCOTT A. CUTHBERT: Augmented. So anything bigger than major is augmented. I want to make sure that this year, we don't lose any points on the problem set from us teaching the computer wrong. So good. Augmented second. So what's this going to be?

[INTERPOSING VOICES]

MICHAEL SCOTT A. CUTHBERT: Augmented augmented second or doubly augmented second. So good. So what do we have to add? What do we--

AUDIENCE: Triply augmented.

MICHAEL SCOTT A. CUTHBERT: Triply augmented second.

AUDIENCE: Quadruple.

MICHAEL SCOTT A. CUTHBERT: Quadruple augmented second. You know those quadruple augmented seconds.

[PLAYING NOTES]

Right? Let's do your tonal theory. Is that consonant or dissonant?

AUDIENCE: Consonant.

AUDIENCE: Consonant.

MICHAEL SCOTT A. CUTHBERT: Consonant. Anybody else? All diminished or augmented intervals or doubly diminished or doubly augmented, whatever are traditionally considered dissonances, No matter what they sound like. So if you doubly augment, triply augmented, quadruply augment, they will generally be regarded as dissonances because they're either a dissonance, or you have written them wrong. And the one that comes up is an augmented second, which sounds just like the beautiful minor third there.

So we have quadruple augmented-- you were very, very close, by the way. You should feel proud on that. Quadruply augmented second. Ascending or descending?

AUDIENCE: Ascending?

MICHAEL SCOTT A. CUTHBERT: Ascending. Good. Great. And Misha, you get this one?

AUDIENCE: I'm not sure.

MICHAEL SCOTT A. CUTHBERT: OK, good. Pass on to Jason?
Good.

AUDIENCE: Doubly diminished-- make doubly diminished second.

MICHAEL SCOTT A. CUTHBERT: Doubly diminished second.
Let's double-check that here. That's a minor second, right?
So diminished second. Doubly diminished second. Great.
Ascending or descending?

AUDIENCE: Spelt this way, it's ascending.

MICHAEL SCOTT A. CUTHBERT: Spelt this way, it's
ascending. And so let's hear those two notes.

[PLAYING NOTES]

Ascending doubly diminished second. Anyone want to make another
argument?

AUDIENCE: [LAUGHS]

MICHAEL SCOTT A. CUTHBERT: Yes, Adam.

AUDIENCE: Just consider what the final pitches are actually.

MICHAEL SCOTT A. CUTHBERT: What the final pitches are.
What do you mean by final pitches?

AUDIENCE: Actually, it turns out to be an F and the E.

MICHAEL SCOTT A. CUTHBERT: Eventually, it turns out to be-
- Adam has an ontology that privileges the ear over the
written ontology. That will do him very well in most
contexts, except for in a sheet music factory, right? Yeah. So
sometimes we want that, sometimes we won't. But yeah, so
you could do that. Any other argument about when? So we
have ascending but sounds ascending or descending sounds
ascending. If you wanted to say I'm on MIDI note, what is
that? 65. This is MIDI note 65. And this one is MIDI note 64.

What's the chromatic number of semitones distance? Sorry for writing so small. Minus 1, right? So if we don't care about spelling, it is a negative interval. Here is a context where you might care about how it's written, in particular. Let's say for some deranged reason, you needed to write this as a chord. Oops, sorry. There are certain rules about where the accidentals go from bottom note to next note. So these are some of the things that come up with computers. These are underdefined parts of music theory.

Underdefined parts of music theory that will come up quite often, that nobody has absolutely decided on these things. Because you will see that they do come up in real music, but they're very, very rare. And I mean, believe the number of doubly diminished octaves in the music21 corpus of tonal music of music of 18th, 19th century. Out of millions of notes, there's about 28 times that it happens. It happens once in the Beach-- no, it happens in the Clara Schumann piece. So it does happen in pretty normal music. But when you see this, what do you think is happening, by the way when you might have this happening at the same time? What's the context? Adam.

AUDIENCE:

The two parts are having a crossover.

MICHAEL SCOTT A. CUTHBERT: Two parts are having a crossover. What's the top part doing? The one that starts off higher?

AUDIENCE:

The top part is going down. It's the one with the right note.

MICHAEL SCOTT A. CUTHBERT: The one with the--

AUDIENCE:

The F.

MICHAEL SCOTT A. CUTHBERT: --F flat. Yeah, so sometimes when you're writing a chromatic scale for an instrument or something, you'll use a lot of flats. And then the other one could be going up. And I think you'll find that this happens in John Philip Sousa's march, *The Thunderer*.

[SCAT SINGING]

It's a melody in the baseline.

[SCAT SINGING]

And that there's-- I think that's one where there's a double-- there's either a diminished octave or a doubly missed octave in the piano part. And in the instrumental parts, you get these doubly diminished seconds and diminished seconds that happen there. Good. So one of the debates is whether or not this number tells you ascending or descending versus the other one. Now, last one. We ended with Jason-- Vincent, did you-- no. Go ahead.

AUDIENCE:

We diminished unison? Diminish one?

MICHAEL SCOTT A. CUTHBERT: Diminish one. Ascending or descending?

AUDIENCE:

Null.

MICHAEL SCOTT A. CUTHBERT: Null.

[LAUGHTER]

Null pointer exception. What do you want to call it then? Let's come up with a music term besides null.

AUDIENCE:

Like a new term?

MICHAEL SCOTT A. CUTHBERT: No, anyone know a music term that's neither ascending or descending for an interval? You don't usually use it in intervals, but you might use it for types of motion. For stationary, we use another one, sometimes oblique. Has anyone heard that term? No. OK, so some people have when-- there's some motion and some not motion. So some people say that. Yeah, so there's kind of a debate. So you could say it's no motion. It's a diminished unison.

The other thing that some people want to do is to make everything always be positive. So you can call this also-- what's another term we could also say? I'm just going to get a fresh thing over here so we can look at that from E-natural to E-flat. We can call this diminished unison. What would we also call this?

AUDIENCE:

Augmented unison.

MICHAEL SCOTT A. CUTHBERT: Augmented unison. And then what's the directionality?

AUDIENCE:

Descending.

MICHAEL SCOTT A. CUTHBERT: Descending. So some people-- descending augmented unisons. Some People who have done some programming things have thought about maybe the diminished unison should vanish from our vocabulary because it's unclear whether music works like integer multiplication, that if you have a descending diminished unison, descending diminished unison, does that mean descending multiplied by diminished unison, which is negative? Negative and negative means you're going up. A negative times a negative is a positive.

At least the descending augmented unison escapes that argument. So some of the weird things that come up. Great. Let's start. Any questions on this before we move on? Thanks. Let's get a head start on one of the hardest parts of problem set 4. What I'd like you to do is given two music21 notes, first, give me what the number 1, 2, 3-- we sometimes call this the generic interval. Some people use the term diatonic interval for this. But we'll use diatonic interval for things like perfect unison, major second, minor third, and so on.

First, given any two notes, what's the generic interval between them? And then let's see if we can work up to what the diatonic interval is. Just as a reminder, there is a class called interval is still locked on this. So we'll do this for about the last six minutes. And that's how we'll end the class today. I'll just be going around and helping people get started. Thanks.